U.S. Department
of Transportation

Research and Innovative Technology
Administration

Volpe National Transportation
Systems Center

A Review of Human-Automation Interaction Failures and Lessons Learned

NASA Airspace Systems Program

Final Report
October 2006

Thomas B. Sheridan and Eric D. Nadler

Notice

This document is disseminated under the sponsorship of the Department of Transportation in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.

Notice

The United States Government does not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to the objective of this report.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE October 2006	3. REPORT TYPE AND DATES COVERED Final Report September 2005 to September 2006

4. TITLE AND SUBTITLE
A Review of Human-Automation Interaction Failures and Lessons Learned

5. FUNDING NUMBERS
NA23/DM345

6. AUTHOR(S)
Thomas B. Sheridan and Eric D. Nadler

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
U.S Department of Transportation
Research and Innovative Technology Administration
John A. Volpe National Transportation Systems Center
55 Broadway, Cambridge, MA 02142-1093

8. PERFORMING ORGANIZATION REPORT NUMBER
DOT-VNTSC-NASA-06-01

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Washington, DC 20546-0001

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
This paper can be accessed at http://www.volpe.dot.gov/hf/pubs.html/.

12a. DISTRIBUTION/AVAILABILITY STATEMENT

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

This report reviews 37 accidents in aviation, other vehicles, process control and other complex systems where human-automation interaction is involved. Implications about causality with respect to design, procedures, management and training are drawn. A number of caveats and recommendations from the salient literature are discussed with respect to human-automation interaction.

14. SUBJECT TERMS
Next Generation Air Transportation System (NGATS), human factors, automation, aviation accidents

15. NUMBER OF PAGES
44

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT Unlimited

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18
298-102

METRIC/ENGLISH CONVERSION FACTORS

ENGLISH TO METRIC			METRIC TO ENGLISH		
LENGTH (APPROXIMATE)			**LENGTH (APPROXIMATE)**		
1 inch (in)	=	2.5 centimeters (cm)	1 millimeter (mm)	=	0.04 inch (in)
1 foot (ft)	=	30 centimeters (cm)	1 centimeter (cm)	=	0.4 inch (in)
1 yard (yd)	=	0.9 meter (m)	1 meter (m)	=	3.3 feet (ft)
1 mile (mi)	=	1.6 kilometers (km)	1 meter (m)	=	1.1 yards (yd)
			1 kilometer (km)	=	0.6 mile (mi)
AREA (APPROXIMATE)			**AREA (APPROXIMATE)**		
1 square inch (sq in, in^2)	=	6.5 square centimeters (cm^2)	1 square centimeter (cm^2)	=	0.16 square inch (sq in, in^2)
1 square foot (sq ft, ft^2)	=	0.09 square meter (m^2)	1 square meter (m^2)	=	1.2 square yards (sq yd, yd^2)
1 square yard (sq yd, yd^2)	=	0.8 square meter (m^2)	1 square kilometer (km^2)	=	0.4 square mile (sq mi, mi^2)
1 square mile (sq mi, mi^2)	=	2.6 square kilometers (km^2)	10,000 square meters (m^2)	=	1 hectare (ha) = 2.5 acres
1 acre = 0.4 hectare (he)	=	4,000 square meters (m^2)			
MASS - WEIGHT (APPROXIMATE)			**MASS - WEIGHT (APPROXIMATE)**		
1 ounce (oz)	=	28 grams (gm)	1 gram (gm)	=	0.036 ounce (oz)
1 pound (lb)	=	0.45 kilogram (kg)	1 kilogram (kg)	=	2.2 pounds (lb)
1 short ton = 2,000 pounds (lb)	=	0.9 tonne (t)	1 tonne (t)	=	1,000 kilograms (kg)
				=	1.1 short tons
VOLUME (APPROXIMATE)			**VOLUME (APPROXIMATE)**		
1 teaspoon (tsp)	=	5 milliliters (ml)	1 milliliter (ml)	=	0.03 fluid ounce (fl oz)
1 tablespoon (tbsp)	=	15 milliliters (ml)	1 liter (l)	=	2.1 pints (pt)
1 fluid ounce (fl oz)	=	30 milliliters (ml)	1 liter (l)	=	1.06 quarts (qt)
1 cup (c)	=	0.24 liter (l)	1 liter (l)	=	0.26 gallon (gal)
1 pint (pt)	=	0.47 liter (l)			
1 quart (qt)	=	0.96 liter (l)			
1 gallon (gal)	=	3.8 liters (l)			
1 cubic foot (cu ft, ft^3)	=	0.03 cubic meter (m^3)	1 cubic meter (m^3)	=	36 cubic feet (cu ft, ft^3)
1 cubic yard (cu yd, yd^3)	=	0.76 cubic meter (m^3)	1 cubic meter (m^3)	=	1.3 cubic yards (cu yd, yd^3)
TEMPERATURE (EXACT)			**TEMPERATURE (EXACT)**		
[(x-32)(5/9)] °F	=	y °C	[(9/5) y + 32] °C	=	x °F

QUICK INCH - CENTIMETER LENGTH CONVERSION

Inches	0			1			2			3			4		5
Centimeters	0	1	2		3	4		5	6		7	8	9	10	11	12	13

QUICK FAHRENHEIT - CELSIUS TEMPERATURE CONVERSION

°F	-40°	-22°	-4°	14°	32°	50°	68°	86°	104°	122°	140°	158°	176°	194°	212°
°C	-40°	-30°	-20°	-10°	0°	10°	20°	30°	40°	50°	60°	70°	80°	90°	100°

For more exact and or other conversion factors, see NIST Miscellaneous Publication 286, Units of Weights and Measures. Price $2.50 SD Catalog No. C13 10286

Updated 6/17/98

TABLE OF CONTENTS

List of Tables

1.0 INTRODUCTION AND SCOPE

The purpose of this review is to consider a variety of failure events where human users interacted with automation, some sophisticated and some not, and to suggest lessons learned from these experiences. Also included are caveats identified in research literature on human-automation interactions that can be applied to design of the Next Generation Air Transportation System (NGATS).

Some events in our sample are failures involving aircraft; others are human interactions with devices in other domains. In almost every case it is not random, unexplainable machine failure or human failure. Rather, it is poor human-machine system design from a human factors perspective: circumstances that are preventable. And while some of these failures have complex causal explanations, most were caused by relatively simple elements of hardware, software, procedure design, or training that were overlooked.

Several accidents not automation-related are included at the end to help make the point that serious consequences can result from simple human user misjudgments in interaction with the physical environment.

Each of the brief summaries was paraphrased from the much longer reports cited. Individual references are listed in parentheses after each heading. These references contain background information and sometimes a colorful description of the unfolding events. Following the failure events summaries are lessons learned and important caveats identified from the literature.

2.0 FAILURE EVENTS INVOLVING AIRCRAFT

2.1 Korean Airlines Flight 007 747 Shot Down by Soviet Air Defense Command (flaw in mode indication)

In August 1983, two minutes after takeoff from Anchorage, pilots engaged the autopilot in “heading” mode and set it directly to the Bethel waypoint. From the black box recordings it appears the inertial navigation system never engaged. This could be because the aircraft was either more than 7.5 miles off the flight route to additional selected waypoints or it was not sufficiently headed in that direction. As a result, the 747 stayed in “inertial navigation armed” mode as the system resorted to the last set “heading” mode as it waited for the required conditions and continued to drift off course.

That early 747 apparently lacked an indicator that the heading mode was the one that was active. (Most likely, the only indication was that the indicator light for the inertial navigation system was amber when it should have been green). The aircraft continued off course and overflew the Soviet Kamchatka Peninsula, which juts into the Bering Sea, then headed straight toward a submarine base. Because of darkness the crew could not see this happening.

MiG fighters were scrambled and chased the 747 for a time, but turned back. By then the aircraft had drifted well off path and soon was over the Soviet territory of Sakhalin Island, where two more MiG fighters were dispatched. They misidentified the aircraft as a U.S. Air Force RC-135, essentially the same as a 747. The Korean aircraft was not on an emergency radio frequency. It was initiating communication with Tokyo, and it did not pick up any Soviet Air Force warning. At that moment Tokyo gave the instruction to climb. This was interpreted by the pursuing Soviet pilot as an evasive maneuver. The MiG pilot was instructed to shoot and did so (Degani, 2004).

2.2 China Airlines 747 Engine Malfunction Near California (over-reliance on autopilot after fatiguing flight)

In February 1985, toward the end of a fatiguing flight from Taipei, the 747-SP lost the rightmost engine and began a right roll due to asymmetric thrust. The autopilot countered by trying to roll left. Since the pilot was hands off in trying to diagnose the cause, he did not notice the only indications of the autopilot effort: the control wheel left rotation, as well as a side slip and a reduction in speed. After some delay the pilot switched the autopilot from FMS to pitch-hold mode but still saw no indication that the autopilot was at its limit in trying to correct the rotation. The aircraft pitched down, the right wing finally dropped, and eventually the pilot switched to manual. The pilot was able to regain control at 9,500 feet (ft) and land safely at San Francisco. The event was attributed to fatigue and boredom at the end of a long flight, forgotten training that indicated manual takeover in such an event, and a lack of instrument indications (Degani, 2004).

2.3 Simmons Airlines ATR-72 Crash Near Chicago (icing disengaged autopilot, surprise manual recovery failed)

In 1994, the ATR-72 encountered icing at 16,000 ft and was instructed to descend and maintain 10,000 and subsequently 8,000 ft. The crew could see the large amount of ice buildup on the wings (more on the right wing than the left). Unknown to the crew, the autopilot was countering a tendency to turn right. Eventually the autopilot reached the limit of its ability and (by design) automatically disengaged. This caused the aircraft to suddenly corkscrew into a sharp right turn, right roll, and 15-degree pitch down. The surprised crew was unable to regain control. Sixteen passengers perished in the crash (Degani, 2004).

2.4 Lockheed L-1011 Crash Over the Florida Everglades (automation state change not communicated to pilot)

In this 1972 incident, the entire flight crew was engaged in troubleshooting a problem with a landing gear indicator light and did not recognize that the altitude hold function of the autopilot had been inadvertently switched off. Meanwhile the aircraft slowly descended into the Florida swamp.

Although several factors contributed to this accident, a major factor was poor feedback on the state of automation provided by the system. The disengagement of automation should have been clearly signaled to the human operator so that it could have been validated. Most current autopilots now provide an aural and/or visual alert when disconnected. The alert remains active for a few seconds or requires a second disconnect command by the pilot before it is silenced. Persistent warnings such as these, especially when they require additional input from the pilot, are intended to decrease the chance of an autopilot disconnect or failure going unnoticed. (National Transportation Safety Board [NTSB], 1973)

2.5 A300 Accident Over the Florida Coast (state transition not communicated to pilot)

Two decades after the above L-1011 accident, an Airbus A300 experienced a similar in-flight incident off the coast of Florida (NTSB, 1998a). At the start of a descent into the terminal area, the autothrottles were holding speed constant, but unknown to the pilots, they were no longer controlling the airspeed when the aircraft leveled off at an intermediate altitude. The aircraft slowed gradually to almost 40 knots (kts) below the last airspeed set by the pilots and stalled after the stall warning activated. There was no evidence of autothrottle malfunction. The crew apparently believed that the automated system was controlling airspeed; in fact it had disengaged. In this aircraft a single press of the disconnect button will disengage the autothrottle control of airspeed. When the system disengages, the green mode annunciator in the primary flight display changes to amber and the illuminated button on the glareshield used to engage the system turns off.

The NTSB (1998a) noted that the change in the annunciators could serve as a warning. However, the passive way in which the displays were formatted did not attract attention. The NTSB also pointed to autothrottle disconnect warning systems in other aircraft that require positive crew action to silence or turn off. These systems incorporate flashing displays and, in some cases, aural alerts that capture the pilot's attention in the case of an inadvertent disconnect. These systems more rigorously adhere to the principle of providing important feedback to the operator about the state of an automated system. Internal transitions between different machine states or modes are sometimes hidden from the user, and as a result the user is unaware of the true state of the machine. This might lead to annoyance or frustration with simple systems, such as VCR/TV controls, where the user fumbles with adjusting the TV while the control is actually in VCR mode. In more complex systems the lack of salient feedback about automation states can lead to catastrophe (Degani, 2004; Norman, 1990).

2.6 A300 Crash in Nagoya (pilot misunderstanding of how automation worked)

In 1994, an A300 crashed in Nagoya, Japan, after the pilots inadvertently engaged the autopilot's go-around mode. The pilots attempted to counter the unexpected pitch-up by making manual inputs, which turned out to be ineffective (Billings, 1997). The pilot attempted to continue the approach by manually deflecting the control column. In all other aircraft, and in this aircraft in all modes except the approach mode, this action would normally disconnect the autopilot. In this particular aircraft, the autopilot has to be manually deselected and cannot be overridden by control column inputs. Consequently, a struggle developed between the pilot and the autopilot, with the pilot attempting to push the nose down through elevator control and the autopilot attempting to lift the nose up through trim control. This caused the aircraft to become so far out of trim that it could no longer be controlled.

These types of misunderstandings result from a mismatch of the pilot's mental model and the behavior of the automated system programmed by the designers (Sherry and Polson, 1999). Several other examples of incidents and accidents resulting from these system misunderstandings have been reported (Billings, 1997; Funk et al., 1999; Sarter and Woods, 1995). While some have had benign outcomes and simply become "lessons learned," others have involved serious loss of life (Leveson, 2004).

2.7 Non-identified General Aviation Crash (pilot impatience, lack of training or judgment)

In 1997, a single-engine airplane operated by a non-instrument-rated pilot took off under instrument meteorological conditions. About two hours later, after following a meandering course, which included reversals and turns of more than 360 degrees, the aircraft crashed into trees at the top of a ridge. No mechanical problems with the airplane's controls, engine, or flight instruments were identified. A person who spoke with the pilot before departure stated that the pilot "... was anxious to get going. He felt he could get above the clouds. His GPS was working and he said as long as he kept the [attitude indicator] steady he'd be all right. He really felt he was going to get above the clouds."

Undoubtedly, many factors played a role in this accident, but the apparent reliance on GPS technology, perhaps to compensate for insufficient training and lack of ratings, stands out as a compelling factor. This general aviation accident further exemplifies the danger of over-reliance on automated systems (NTSB, 1998b).

2.8 American Airlines B-757 Crash Over Cali, Columbia (confusion over FMS waypoint codes)

Two significant events in the loss of a B-757 near Cali, Colombia, in 1995, were the pilot asking for clearance to take the Rozo approach followed by the pilot typing “R” into the FMS. The pilot should have typed the four letters “ROZO” instead of “R.” The latter was the symbol for a different radio beacon (called Romeo) near Bogota. As a result, the aircraft incorrectly turned toward mountainous terrain.

While these events are non-controversial, the link between the two events could be explained by any of the following (Leveson, 2001):

- *Crew Procedure Error*: In the rush to start the descent, the captain entered the name of the waypoint without normal verification from the other pilot.
- *Pilot Error*: In the rush to start the descent, the pilot executed a change of course without verifying its effect on the flight path.
- *Approach Chart and FMS Inconsistencies*: The identifier used to identify ROZO on the approach chart (*R*) did not match the identifier used to call up ROZO in the FMS.
- *FMS Design Deficiency*: The FMS did not provide the pilot with feedback that choosing the first identifier listed on the display was not the closest beacon with that identifier.
- *American Airlines Training Deficiency*: The pilots flying into South America were not warned about duplicate beacon identifiers and were not adequately trained on the logic and priorities used in the FMS on the aircraft.
- *Manufacturers’ Deficiencies*: Jeppesen-Sanderson did not inform airlines operating FMS-equipped aircraft of the differences between navigation information provided by Jeppesen-Sanderson FMS navigation databases, Jeppesen-Sanderson approach charts, or the logic and priorities used in the display of electronic FMS navigation information.
- *International Standards Deficiency*: There was no single worldwide standard for the providers of electronic navigation databases used in flight management systems.

In addition to the pilot not starting with an accurate mental model, a mental model may later become incorrect due to lack of feedback, inaccurate feedback, or inadequate processing of the feedback. A contributing factor cited in the Cali B-757 accident report was the omission of the waypoints behind the aircraft from cockpit displays, which contributed to the crew not realizing that the waypoint they were searching for was behind them (missing feedback) (Leveson, 2004).

2.9 A320 Crash in Bangalore, India (control mode error, misunderstanding the automation)

In this accident the pilot had disconnected his flight director during approach and assumed that the co-pilot would do the same. The result would have been a mode configuration in which airspeed was automatically controlled by the autothrottle (the speed mode), which is the recommended procedure for the approach phase. Since the co-pilot had not turned off his flight director, the open descent mode activated when a lower altitude was selected instead of speed mode, eventually contributing to the crash of the aircraft short of the runway (Sarter and Woods, 1995).

2.10 Aero Peru 613 Crash (pitot tubes taped for painting: sloppy maintenance, poor inspection by pilot)

After a routine walk-around pilot inspection of a freshly painted Aero Peru B757, the aircraft took off for a night IFR flight from Peru to Chile. Immediately, the pilot could not understand what was happening to the aircraft. Neither the altimeter nor the airspeed indicators made any sense, with readings as though the aircraft was still on the ground. A request was made to air traffic control (ATC), which gave them their speed and altitude, but these numbers disagreed with the instruments. Alarms started to go off. Forty miles over the Pacific, they requested vectors for circling and returning to the Lima airport. Soon the crew requested another aircraft rendezvous with them and guided them back. Lima dispatched a 707. The crew was confused when an overspeed alarm sounded at the same time as a stall warning, with no correlation to throttle settings. After receiving a ground proximity warning and trying to pull up, they felt the aircraft bounce off the water. The resulting damage was too great and the aircraft rolled over. The aircraft crashed into the Pacific, killing all aboard.

The investigation found that during the painting all the pitot tubes were covered with plastic tape and the tape was never removed. The pilots had not noticed this in the walk-around. ATC was giving them speed and altitude based on the transponder, which was incorrect (Casey, 2006).

2.11 2002 Midair Collision Over Uerberlingen, Germany (pilot decision to follow ATM advice rather than TCAS resolution advisory)

On the night of July 1, 2002, a DHL B757 collided with a Bashkirian Airlines Tupolev-154 over Lake Constance in Germany, resulting in 71 fatalities. After contacting the 757, a Swiss controller (Zurich Center, by agreement with the German government) issued two successive clearances to climb. There was no further contact with either aircraft until the TCAS gave both pilots a traffic advisory. At that point the controller instructed the T154 to descend. The TCAS then instructed the T154 to climb and the 757 to descend. The T154 pilot chose to obey the controller rather than TCAS and descended, while the 757 did as instructed by TCAS and also descended. They collided at FL 350.

A number of factors contributed to the accident. First, the pilot of the Tupolev, whose command of English was questionable, was apparently slow in responding to the controller's descent instructions before the TCAS issued the contrary resolution (which had seemed warranted due to low expected traffic). Second, even though the controller was also at fault for not detecting the conflict earlier, there was only a single controller working, and he was monitoring two different radars. Third, the radar had been downgraded as the main radar system was out for maintenance and a backup was in use. Fourth, German controllers working the Karlsruhe sector had noticed the unfolding situation and had tried to contact the Swiss controller on the telephone, but the telephone was non-functional due to lack of maintenance and the German controllers could not get through. Finally, the TCAS communicates only with the pilot; it does not communicate its resolution instructions to the controller, and the Swiss controller had no way of knowing the TCAS was giving instructions that were in conflict with his own. As a result, the controller did not anticipate a conflict.

This accident is an example of Reason's (1990) "Swiss cheese" model, i.e., when multiple risk factors come together accidents are likely to happen. To add to the original tragedy, the controller was later murdered by a man whose wife and children died in the crash, and eight employees of the Swiss Skyguide navigation service were charged with manslaughter for "organizational shortcomings" (Nunes and Laursen, 2004).

2.12 2004 Roller Coaster Ride of Malaysia Airlines B777 (unanticipated software failure)

On a flight from Australia to Malaysia, the aircraft suddenly climbed 3,000 ft. The pilot immediately disconnected the autopilot and pointed the nose down. The plane then jerked into a steep dive. The pilot throttled back and tried to slow the aircraft, but it again raced into another climb. The pilot finally regained control.

An investigation revealed a defective software program that provided incorrect speed and acceleration data. This confused the flight computers and ignored some of the pilot's commands. Boeing mandated a software fix. While no crash has yet been attributed to serious software errors, there are concerns that it surely will happen (Michaels and Pasztor, 2006).

2.13 October 2005 British Airways A319 Electronics Failure (unanticipated and unreplicated software problem)

On a night flight from London to Budapest, nearly all cockpit lights and electronic displays, along with radios and autopilot systems, went dark for 90 seconds. Efforts to replicate the failure failed. (Allegedly, British Airways had already discovered four similar cases on Airbus jets) (Michaels and Pasztor, 2006).

2.14 Embraer Test Flight: One-Minute Blackout of Computer Displays (presumably due to a software glitch)

During an Empresa Brasileira de Aeronautica SA test flight of a new jetliner, the cockpit displays went dark for one minute, then came back on. Embraer says it has fixed the problem. FAA has since ordered programming changes to the plane (Michaels and Pasztor, 2006).

2.15 2003 Crash of Air Midwest/U.S. Airways Express Beech 1900D (shortcutting of required maintenance procedures)

During a routine maintenance job at night, a third-party subcontracted maintenance trainee was tasked to check the tensioning of control cables. The paperwork called for an apparently complex procedure to access and retension cables, including removing seats and other parts. Some cables appeared to need retensioning, in particular the elevator cables. The trainee and his supervisor discussed the task and decided that the procedure was more complex than necessary and in any case was probably for a different-model aircraft. They decided the cables could be tensioned in a simpler manner and shortcut many of the recommended maintenance procedures for tensioning them (and in the process shortening them).

The first flight following the tensioning maintenance had a full load of 19 passengers and a larger than normal batch of luggage. Immediately upon takeoff, after the wheels came up, the captain noticed that it was hard for the nose to go down: full forward yoke reached a limit that was not enough, and the nose stayed up. When the nose pointed 68 degrees upward and the aircraft had slowed to 31 knots (kts), the aircraft stalled. Still over the runway, it rolled, fell, and crashed, killing all aboard.

Three factors contributed to the accident. First, the aircraft was unusually tail-heavy with passengers and luggage. Second, an error was made in computing the excess weight of passengers and luggage (the conventions used in the computations for passengers and luggage were too light). Finally, the combination required full forward yoke, but the cable hit-stop and elevator nose-down had been adjusted 2 inches short of full forward nose down position as part of the improvised maintenance shortcut (Casey, 2006).

2.16 John Denver Crash into the Pacific (cutting corners in manufacture, poor human interface)

John Denver, an experienced pilot and popular singer, bought a Burt Rutan, designed and “homebuilt” Long EZ experimental aircraft (pusher propeller, canard). To simplify fuel piping, in contrast to the original Rutan design the fuel valve (to switch fuel tanks) had been placed in the rear passenger compartment, reachable only by releasing the control stick, turning around, and using a short pole to operate the valve in an awkward direction. The float-ball indicator of available fuel was also nonlinear, showing one-third full when the tank was actually almost empty.

Denver trusted his maintenance technician who, not aware of this nonlinearity, said there was quite enough fuel for a few touch-and-goes that Denver wanted to practice at the Monterey Peninsula Airport in California. When the engine surprised him by running out of fuel, Denver apparently tried to turn and operate the fuel valve but inadvertently activated a pedal (rudder control). This stalled the aircraft, which rolled and crashed into Monterey Bay (Casey, 2006).

2.17 U.S. Soldier in Afghanistan Inadvertently Calls for Air Strike on Own Position (ignorance of reset operation)

In December 2001, a U.S. soldier in Afghanistan needed to call in an air strike using a handheld GPS /communication device. After entering the coordinates for the strike into the device, he noticed that it indicated a need for battery replacement. Having a spare battery available and thinking that he had better do that quickly, he replaced the battery. What the soldier did not realize was that when the battery was replaced the system reset to the present (his own) coordinates. The air strike came down on his position, killing two-dozen people (Casey, 2006).

2.18 Loss of Black Hawk Helicopters to Friendly Fire (ill-defined procedures and traffic management responsibilities)

A major factor in the loss of the Black Hawk helicopters to friendly fire over northern Iraq is that they normally flew only in the boundary areas of the No Fly Zone, and procedures for handling aircraft in those areas were ill-defined (Leveson, Allen, and Storey, 2002). An Army base controlled the flights of the Black Hawks while an Air Force base controlled all other airspace activity. A common control level was higher in the control management structure than the level at which the decisions were made in this accident. Communication problems existed between the Army and Air Force bases at the intermediate control levels. The aircraft surveillance officer (ASO) thought she was responsible *only* for identifying and tracking aircraft south of the 36th parallel. The air traffic controller for the area north of the 36th parallel thought the ASO was also tracking and identifying aircraft in his area and accordingly took no action (Leveson, 2004).

2.19 Upset in Descent of NASA M2F2 Lifting Body (design led to pilot control reversal)

An experimental wingless aircraft was launched by falling free from a B52. In its initial flight in July 1966, the test pilot tried to adjust the ratio control sensitivity. Unfortunately, he committed a control reversal that made the stick more sensitive rather than less. This resulted in a severe "pilot-induced oscillation," predictable for an aircraft of this configuration. Almost too late, he desperately let go of the stick and the aircraft finally stabilized. He then realized what he had done (Casey, 2006).

2.20 Concorde Crash Precipitated by Runway Debris (control tower automation may reduce controller vigilance of airport surface)

The Air France Concorde, a supersonic passenger aircraft, crashed in July 2000 while attempting to depart from Paris. The crash killed all 109 on board and four people on the ground. Even though automation did not contribute to this particular crash, it is included in this paper because of current trends toward automation in air traffic control towers (e.g., "virtual" towers located away from the runways and depending on instrument surveillance). These virtual towers have the potential to divert attention from scanning the airport surface for debris, as controllers currently do today, while updating their awareness of aircraft position and movements. A landing gear tire on the Concorde encountered a 17-by-1-inch metal strip that had fallen from a DC-10 five minutes earlier.

The Concorde had a history of tire bursts and deflations 60 times higher than subsonic jets (e.g., one in 3,000 flights compared to one in 100,000 flights for the A340). Heavy, fast-moving pieces of the tire severed and shorted electrical wiring in the wing, igniting fuel adjacent to the wheel well, and a fire erupted when the fuel tank burst. The flight engineer idled the second left engine 12 seconds after its fire alarm began. The Concorde had gained less than 400 ft in altitude when the flight crew lost control, and the aircraft crashed four miles from the end of the runway. The Bureau Enquetes-Accidents (BEA), the French equivalent of the NTSB, attributed the crash in part to the performance of only two runway examinations at Charles de Gaulle International Airport, when three were specified in a service memo. However, the additional daily inspection may not have discovered the metal debris during the five minutes between its appearance and the Concorde departure. British Airways terminated Concorde service in October 2003 for reasons of profitability (*Air Safety Week,* 2002).

3.0 FAILURE EVENTS IN OTHER TRANSPORTATION SYSTEMS

3.1 *Royal Majesty* Grounding (over-reliance on automation, lack of failure awareness)

This example from the maritime industry illustrates the effects of over-reliance on automated systems.

The cruise ship *Royal Majesty* ran aground off Nantucket after veering several miles off course toward shallow waters. Fortunately, there were no injuries or fatalities as a result of the accident, but losses totaled $2 million in structural damage and $5 million in lost revenue. The automated systems in this ship included an autopilot and an automatic radar plotting aid that was tied to signals received by a GPS. Under normal operating conditions, the autopilot used GPS signals to keep the ship on its intended course. However, the GPS signals were lost when the cable from the antenna frayed (it was placed in an area of the ship where many sailors walked). As a result, the GPS and autopilot automatically and without warning switched to dead-reckoning mode, no longer correcting for winds and tides, which carried the ship toward the shore.

According to the NTSB report on the accident, the probable cause was the crew's over-reliance on the automatic radar plotting aid and management's failure to ensure that the crew was adequately trained in understanding the automation features, capabilities, and limitations. The report went on to state that "the watch officers' monitoring of the status of the vessel's GPS was deficient throughout the voyage …" and that "all the watch-standing officers were overly reliant on the automated position display and were, for all intents and purposes, sailing the map display instead of using navigation aids or lookout information."

This accident represents a classic case of automation complacency related to inappropriately high trust in the automation. It also demonstrates the importance of salient feedback about automation states and actions. The text annunciators that distinguished between the dead-reckoning and satellite modes were not salient enough to draw the crew's attention to the problem (Degani, 2004; Lee & See, 2004; see Degani for a more detailed account of the accident).

3.2 *Herald of Free Enterprise* Sinking off Zeebrugge, Netherlands (poor management planning)

In March of 1987, the roll-on, roll-off ferry *Herald of Free Enterprise* was en route to Dover with her bow doors open. Shortly after departure, water came over the bow sill and flooded the lower deck. She sank in less than 2 minutes, drowning 150 passengers and 38 crew. The accident investigation pointed to lax management, both on board and on shore, and the crew's lack of comprehension of their duties (Reason, 1990).

3.3 BMW 7 Series iDrive Electronic Dashboard (designer gadget fantasy gone wild)

The 2003 BMW 7 series featured an electronic dashboard called “iDrive” that had roughly 700 features and was satirized by the automotive press. *Car and Driver* called it “a lunatic attempt to replace intuitive controls with silicon, an electronic paper clip on a lease plan.” *Road and Track* headlined an article “iDrive? No, you drive, while I fiddle with the controller” and asserted that the system “forced the driver to think too much” (just the opposite of good human factors engineering) (Vicente, 2004).

3.4 Milstar Satellite Loss (poor assumptions and lack of design coordination)

A Milstar satellite was lost due to inadequate attitude control of the Titan/Centaur launch vehicle, which used an incorrect process model based on erroneous inputs in a software load tape. After the accident, it was discovered that no one had tested the software using the actual load tape and that all software testers had assumed someone else was doing so. (Leveson, 2001) System engineering and mission assurance activities were missing or ineffective, and individual development and assurance groups did not have a common control or management function (Leveson, 2004).

3.5 Failed Ariane Liftoff (poor assumptions in anticipating of software requirement)

Even though the French Ariane 5 spacecraft trajectory had been changed from that of the Ariane 4, the inertial reference system software had not been updated sufficiently, resulting in a failed launch (Leveson, 2004).

3.6 Solar Heliospheric Observatory (failure to communicate a procedure change to operators)

A principal factor in the loss of contact with SOHO (SOlar Heliosperic Observatory) in 1998 was the failure to communicate to operators that a functional change had been made in a procedure to perform gyro spin down (Leveson, 2004).

4.0 FAILURE EVENTS IN PROCESS CONTROL SYSTEMS

4.1 Bhopal, India, Union Carbide Leak (multiple failures in design, maintenance, and management)

From a fatality standpoint, the worst single human-automation accident in history was the Union Carbide plant accident in Bhopal, India, in December of 1984. At least 2,500 people were killed and 200,000 injured by a gas leak involving the influx of water into a storage tank of methyl isocyanate.

The accident was variously attributed to “botched maintenance, operator errors, need for improved by-pass pipes, failed safety systems, incompetent management, drought, agricultural economics and bad government decisions.” The plant should not have been located close to a densely populated area. There were poor evacuation measures, few available gas masks, and an inadequate siren alarm system. The inexperienced operators neglected inspector warnings and failed to notice or control pressure buildup. Pressure and temperature gauges were faulty, there was no indication of valve settings, and scrubbers lacked sufficient capacity. This accident clearly resulted from a concomitance of factors (Reason, 1990).

4.2 Nuclear Meltdown at Three Mile Island (failures in design, procedures, management [including maintenance], training, and regulation)

In March 1979, a turbine tripped at one of the two nuclear plants on Three Mile Island near Harrisburg PA. As a maintenance crew was working on water treatment, some water leaked through a faulty seal and entered the plant’s instrument air system, interrupting the air pressure applied to two feedwater pumps and giving a false indication that something was operating incorrectly. This automatically caused the pumps to stop, cutting water flow to the steam generator, tripping the turbine, and activating emergency feedwater pumps that drew water from emergency tanks. The pipes from these tanks had erroneously been left blocked two days earlier during maintenance. With no heat removal by the cooling water, the temperature rose rapidly and caused an automatic reactor “scram” (boron rods dropped into the reactor core to stop the chain reaction). Pressure from residual heat at that point (now only 13 seconds into the accident) was dissipated through a pressure-operated relief valve that is supposed to open briefly and close. The valve stuck open, and the radioactive water under high pressure poured into the containment area and down into the basement.

The indicator light on the valve showed that it had been commanded to shut, but there was no indication of the actual valve status. The hundreds of alarms that lit up were not organized in a logical fashion and gave little indication of causality. The crew cut back on high-pressure emergency water injection, having been taught never to fill to the limit and not realizing that water was emptying out.

Although no one was killed, loss of cooling water caused significant damage to the reactor core and brought nuclear power plant construction in the U.S. to a halt (Reason, 1990).

4.3 Failure in British Chemical Plant (poor anticipation of unsafe interactions during design)

In a batch chemical reactor in England, a computer controlled the flow of catalyst into the reactor and the flow of water into the reflux condenser to cool the reaction. Sensor inputs to the computer were to warn of any problems in various parts of the plant. The programmers were told that if a fault occurred in the plant, the computer was to leave all controlled variables as they were and sound an alarm. On one occasion, the computer received a signal indicating a low oil level in a gearbox. The computer reacted as its requirements specified. It sounded an alarm and left the controls as they were. By coincidence, a catalyst had been added to the reactor, but the computer had just started to increase the cooling-water flow to the reflux condenser. The flow was therefore kept at a low rate. The reactor overheated, the relief valve lifted, and the contents of the reactor were discharged into the atmosphere.

There were no component failures involved in this accident. Individual components, including the software, worked as specified, but together they created a hazardous system state. Merely increasing the reliability of the components or protecting against their failure would not have prevented the loss. Prevention required identifying and eliminating or mitigating unsafe interactions among the system components (Kletz, 1982; Leveson, 2004).

4.4 Uncontrolled Chain Reaction at Japanese Breeder Reactor (operators' shortcut of recommended safety procedures)

An essential step in enriching uranium fuel involves mixing concentrated uranium powder with nitric acid, a process using a very specialized mixing apparatus that the Japanese Science and Technology Agency had declared involved “no possibility of critical accident occurrence due to malfunction and other failures.” However, for efficiency, local staff had developed shortcut timesaving procedures approved by the manufacturing quality assurance people but not by safety management. Instead of using a tall and narrow extraction and buffer column designed to prohibit accumulation of sufficient mass to cause a chain reaction, they adopted a bowl-like vessel that put the contents together in one big yellow frothy container. Unfortunately, within seconds of two operators dumping the final liquid into the container, an intense blue light shone from the center of the mass; this was a chain reaction. Intense heat ensued, and radiation alarms were set off. It took 20 hours to stop the chain reaction. The two operators were rushed to the hospital with intense radiation burns, but both died within a few days. The company’s fuel reprocessing license was suspended (Casey, 2006).

4.5 Observed Dysfunction in Steel Plant Blast Furnace Department (poor communication regarding authority)

At an iron and steel plant, frequent accidents occurred at the boundary of the blast furnace department and the transport department. One conflict arose when a signal informing transport workers of the state of the blast furnace did not work and was not repaired because each department was waiting for the other to fix it. Such dysfunction was the result of too many management levels separating workers in the two departments from a common manager: the greater the distance, the more difficult the communication, and the greater the uncertainty and risk.

There was also evidence that accidents were more likely in boundary areas or in overlap areas where two or more controllers (human and/or automated) controlled the same process. In both boundary and overlap areas, the potential existed for ambiguity and for conflicts between independently made decisions. When controlling in boundary areas, there was confusion over who was actually in control (which control loop was currently exercising control over the process), leading to missing control actions. The functions in the boundary areas were often poorly defined (Leplat, 1987).

5.0 FAILURE EVENTS IN OTHER SYSTEMS

5.1 The Florida Butterfly Ballot (poor interface design, lack of usability testing)

The infamous “butterfly ballot” used in Palm Beach, Florida, during the 2000 presidential election was laid out and printed in such a way that when people attempted to vote for the Democratic candidate, Al Gore, many ended up voting for the Reform Party candidate, Buchanan, or for both Gore and Buchanan. Although the Democratic candidates were listed in the second item in a left-hand column, the supposedly corresponding hole to punch was the third one down in a centrally located column, where the proper holes alternated between candidates listed in the left-hand column and the right-hand column. This egregious design flaw was a simple matter of very poor human factors in laying out the ballot to accommodate a slightly simpler mechanism rather than implementing what would be simple, obvious, and natural for the human user (Vicente, 2004).

5.2 Emergency MRI Oxygen Bottle Kills Child (lack of anticipation of critical safety requirements)

A child patient undergoing an MRI in Westchester Medical Center in New York needed emergency oxygen. The available supply (from aluminum bottles in the MRI facility) had been exhausted and additional oxygen was requested. A passing nurse heard the call and rushed in with a conventional steel oxygen bottle, which she should have known was inappropriate when used in areas where there are powerful magnets, such as those used in MRI. When the steel bottle was close enough to the magnet, it was suddenly pulled from the nurse’s hands by the magnet and into the MRI cavity. It struck the child’s head and immediately killed him. In the rush to provide the child with oxygen, the nurse did not think about how ferrous material is attracted by magnetism. This accident set off an immediate nationwide effort to improve safety conditions and procedures for MRI facilities (Casey, 2006).

5.3 Production of New Salk Vaccine at Cutter Labs (rush to scale up production precluded precautionary care)

In April 1955, experimental evaluation results of the Salk vaccine were announced. Against the most virulent type of polio, the vaccine proved to be at least 60 percent effective, and against the more common types of polio it was more than 90 percent effective. Six pharmaceutical firms were immediately licensed to produce the vaccine, including Cutter Laboratories of Berkeley, California. The challenge was how to scale the laboratory methods for manufacturing. This included collecting and handling thousands of monkey kidneys and batch processing to deactivate live polio virus, then testing the vaccine. Cutter decided to test random batches due to costs and the urgency in making the product available. Government inspectors verified the paperwork but did not verify the product. Shortly after 308,000 first- and second-graders and 82,000 patients in medical offices were inoculated, there were problems. Two hundred and forty cases of full-fledged polio were reported; 180 of those infected were paralyzed and 11 died. The U.S. surgeon general cancelled the entire program. An investigation at Cutter pointed to “small changes" in the lab procedures originally used to kill the virus to those used in industrial-scale production. Cutter subsequently went out of business (Casey, 2006).

5.4 Patient Morphine Overdose from Infusion Pump (nurses' complaints about programming disregarded)

Patient-controlled analgesia (PCA) infusion pumps are common in hospitals, and many firms manufacture them. Nurses responsible for programming the pumps have long reported that the task is not easy, and several organizations have reported on the potential safety problems. The Emergency Care Research Institute issued an alert about PCAs being "susceptible to mis-programming" and stated that "the user interface and logic of the pump are particularly complex and tedious." In February of 2000, Danielle McCray, who had just given birth to a healthy baby via C-section and was in good health, died of a morphine overdose from a PCA infusion pump. She had received four times the lethal dose. Subsequent investigations estimated that between 65 and 650 deaths have occurred from PCA programming errors. PCAs have a 75 percent market penetration in the U.S. Allegedly, the manufacturer told a reporter that the device "has no design flaws … the pump is safe if used as directed." Litigation resulting from the McCray death pointed to the nurse and hospital rather than the device manufacturer. This, unfortunately, is a typical response: errors in using automation are attributed to "bad apple" users of machines rather than to the design of the machines (Vicente, 2004).

5.5 Olympic Swim Meet Scoring Device that Could Not Be Altered (lack of flexibility in design and systems management)

At the Barcelona Olympics in 1992, the Brazilian swim meet judge Da Silvieri Lobo meant to give champion Canadian diver Sylie Frechette a 9.7 score. Inadvertently, she hit the wrong button on her handheld computer terminal and the score came up 8.7. She tried to make a change by re-entering the score but the system software would not allow the change. Confusion and delay followed, coupled by the difficulty of the Japanese assistant referee in understanding the judge's Portuguese-accented English. The audience demanded a score, and the referees finally decided that the 8.7 would stand (Casey, 2006).

5.6 Counting of Instruments and Sponges in Complex Surgeries (lack of appreciation for workload/distraction effects)

Two surgeons and two human factors professionals (one of whom was the writer) observed 10 complex (up to 10-hour) surgeries in a major Boston hospital. On multiple occasions it was evident that unrealistic expectations were being placed on operating room (OR) scrub and circulating nurses to perform accurate instrument and sponge counts to ensure that what went into the patient also came out. ("Foreign bodies" left inside patients after they have been sewn up cause infection and incur large malpractice damage claims for doctors and hospitals.) In addition to counting, the nurses must do many tasks to assist the surgeons and anesthesiologists. In this operation, simultaneous counting and performing of other duties resulted in counting errors that caused significant delays and the need to x-ray patients when the procedures were almost completed.

The observation study also revealed the lack of sufficient information in handoffs between doctors, between nurses (who may leave because of a shift change in the middle of a procedure), and between doctors and nurses. The counting and handoff problems are now recognized as areas that need to be addressed. Automatic optical scanning of bar-coded sponges and computer-pattern recognition are being developed to make the counting process more reliable and less demanding on the nurses. It remains to be seen whether it will work (Author, personal experience).

5.7 VCR Remote Control (technology overkill)

VCR remote-control devices, or “clickers,” come in many varieties, mostly different from one another, and most seem to have buttons the user never learns to use. The VCR clicker is often cited as an example of technology overkill. Its complexity leads many users to abandon efforts to master it. Fortunately, errors in its use are not life-threatening. For additional information, Degani (2004) has a whole chapter on VCR controls.

6.0 LESSONS LEARNED FROM HUMAN AUTOMATION FAILURES

What are the lessons learned from the 38 human-automation interaction failures in the preceding chapters? Table 1 lists specific reasons for the 38 failure events along with their relevance (in this author’s judgment) to four main causal categories: (1) design of hardware and software, (2) procedure, (3) management, and (4) training.

Note that all categories are well populated. This is not surprising. The four causal factors are interdependent. Interface design and procedures go together. Management is responsible for creating a culture of safety and ensuring that the design is working and the procedures are being followed. Operators will not understand the design and the procedures without proper training. No amount of training can make up for poor design and procedures. If hardware and software are not well designed to make operation efficient and natural, especially when off-normal and rare events occur, then operator error can be expected. Good procedures can enhance efficiency and safety, but with sufficiently off-normal events there may be no established procedures. The operator must figure things out and decide what to do, and management is responsible for selecting operating personnel capable of coping with unexpected events.

Table 1. Judged Reasons for Failure in Events Cited

FAILURES CITED	DESIGN	PROCEDURE	MANAGEMENT	TRAINING
AIRCRAFT				
Flaw in mode indication	X			
Over-reliance on autopilot after fatiguing flight		X		X
Icing disengaged autopilot, manual recovery failed	X			X
Automation state change not communicated to pilot	X		X	
State transition not communicated to pilot	X			
Pilot misunderstanding of how automation worked				X
Pilot impatience, lack or training or judgment				X
Confusion over FMS waypoint codes		X		X
Control mode errors, misunderstanding the automation	X			X
Taped pitot tubes: poor maintenance & inspection by pilot		X	X	X
Pilot failed to follow TCAS advisory		X	X	X
Software bug caused roller-coaster ride	X			
Software bug caused failure of systems and displays	X			
Software bug case blackout of displays	X			
Shortcutting of required maintenance procedures		X	X	
Cutting corners in manufacture; poor human interface	X		X	
Ignorance of reset operation	X			X
Ill-defined procedures and traffic management		X	X	
Poor design led to pilot control reversal	X			X
Control tower automation may reduce runway vigilance			X	
OTHER VEHICLES				
Over-reliance on automation; lack of failure awareness			X	X
Poor management planning	X		X	
Designer gadget fantasy gone wild	X		X	
Poor assumptions and lack of coordination in design	X		X	
Poor assumptions in anticipating software requirement	X		X	
Failure to communicate to operators a procedure change			X	X
PROCESS CONTROL				
Multiple failures in design, maintenance and management	X		X	
Design, maintenance, procedures, management, training	X	X	X	X
Poor anticipation of unsafe interactions during design	X			
Operators shortcut of recommended safety procedures		X	X	
Poor communications regarding authority			X	X
OTHER SYSTEMS				
Poor interface design; lack of usability testing	X		X	
Lack of anticipation of critical safety requirements		X	X	X
Rush to manufacture precluded precautionary care	X	X	X	
Nurses complaints about programming disregarded	X		X	
Lack of flexibility in systems management	X		X	
Lack of appreciation for workload/distraction effects		X	X	X
Technology overkill	X		X	

6.1 Degani's Summary Observations: Another Set of Causal and Remedial Considerations

Degani (2004) concluded his book with a series of 35 observations about human-automation interactions, which are abbreviated, modified, and combined here into 27 observations:

1. Time delays in feedback from the automation confuse and frustrate the user.
2. In many control systems there is an underlying structure where the same user action leads to different end states (or modes). If you don't know the current state you may not be able to predict the future state and can be confused.
3. Both the control interface and the user manual are always highly abstracted descriptions of the machine behavior, so it is imperative to make them correct and clear.
4. Insofar as feasible the same language should be used to describe the machine behavior and the operating instructions to the user (on the user interface and in the manual).
5. Important machine state transitions are typically triggered automatically by physical variables with no input from the user. When these events are not easily identifiable, warnings or other feedback should be used.
6. The onset of a potential problem occurs when what is expected diverges from the reality.
7. Small human errors of forgetfulness or inattention can be magnified because of a powerful or persistent machine.
8. Learned and habituated population stereotypes form an "internal model" that shapes how we users can be expected to respond. Designs should respect these stereotypes rather than violate them.
9. User interaction should be consistent across menus, submenus, and all forms of interaction.
10. Users should be aware that control systems have operating modes (contingencies like alarm on), reference values, or set points which they are trying to track, energy resources that they draw on to manifest their control outputs, disturbances that act to offset their efforts, and control laws that determine the control action given the error or model-based estimation of current state.
11. Just because a mode is nominally switched on does not necessary mean that control is active in that mode.
12. When the user is monitoring and supervising automation but is physically and cognitively separated, the effects of corrective user action are sometimes obscure.
13. The user should be aware of default conditions and settings. The default logic is "culturally embedded," so the designer's default logic may not be the same as the default logic of the user (Infield and Corker, 1997).
14. Proper operation of control modes that are rarely used (e.g., in emergencies) must be readily apparent.
15. While magic acts conceal important information and provide irrelevant and extraneous information, interface design does just the opposite.
16. A minimal number of models, display indications, steps, and events that do the job are usually best.

17. Human error is inevitable. Emphasis should be placed on means for quick recovery.
18. Paths to unsafe regions of operating state space need be identified and be blocked by interlocks, guards, warnings, and interface indications.
19. Execution of correct procedures in time-critical situations can be improved by decision aids and a well-designed interface.
20. A procedure that serves as a band-aid for a faulty and non-robust design may go unused.
21. A change in a reference value can trigger a mode change and vice versa.
22. Users come to trust and depend upon highly reliable controls without always understanding their limitations, which can lead them to ignore and misinterpret cues that the automation is behaving abnormally.
23. Envelope protection (e.g., where a special control mode automatically takes over as an aircraft approaches a stall) is complex to design and should not be taken for granted as providing full safety.
24. Computers and automation systems cannot "think" beyond the state transition rules that have been programmed by designers.
25. The decision to disengage the automation in times of emergency has sometimes led to disaster because the manual recovery effort was inappropriate or not sufficiently fast, but not disengaging has also led to disaster because the automation was not sufficiently robust.
26. Criteria for software verification include error states (when a display is no longer a correct indication of the machine state), augmenting states (when the display or other information says a mode is available when it is not), and restricting states (when the user has no information that certain events can trigger mode or state changes).
27. There are sometimes conditions that drive a system into an unsafe region of state space from which it cannot recover or is difficult to recover. Such conditions can be anticipated in design and measures taken to avoid them.

6.2 Awareness of the Problems

There is a growing literature on human-automation interaction in aviation, both real-world failures such as those described above and laboratory experiments (Wiener and Nagel, 1988; Sheridan (1992, 2002); Wickens et al, 1998; Decker and Hollnagel, 1999; Sarter and Amalberti, 2000; Sheridan and Parasuraman, 2006). It is clear that whatever the domain, the hardware and software are becoming more reliable with time and the problems point increasingly to the human interaction. Perrow (1984), for example, asserts that 60 to 80 percent of accidents are attributed to human error. It is not clear that the (probably unstoppable) trend toward further automation will change this.

By itself, automation (artificial sensors, computer logic, and mechanical actuators combined into control loops to perform given tasks) is not a bad thing. One can argue that it makes life better in numerous ways. The root problem lies in thinking that automation simply replaces people, and that since people are the ones who make errors, there will be fewer system failures when people "are removed from the system." The fact is that people are not removed. Automating simply changes the role of the human user from that of direct, hands-on interaction with the vehicle, process, or device being controlled to that of a supervisor. A supervisor is required to plan the action, teach (program) the computer, monitor the action of the automation, intervene to replan and reprogram either if the automation fails or if it is insufficiently robust, and to learn from experience. (Sheridan, 1992, 2002)

Bainbridge (1987) was among the first to articulate what she called the "ironies of automation." A first irony is that errors by the automation designers themselves make a significant contribution to human-automation failures. A second irony is that the same designer who seeks to eliminate human beings still relies on the human to perform the tasks the designer does not know how to automate.

In reference to the automation trend, Reason (1990) commented that "If a group of human factor specialists sat down with malign intent of conceiving an activity that was wholly ill-matched to the strengths and weaknesses of human cognition, they might well have come up with something not altogether different from what is currently demanded …".

6.3 Function Allocation

Which functions to allocate to humans and which functions to allocate to machines is an old question. Fitts (1951) published what has come to be called his MABA-MABA list:

> *Men (sic) are better at:* detecting small amounts of visual, auditory or chemical energy; perceiving patterns of light or sound; improvising and using flexible procedures; storing information for long periods of time and recalling appropriate parts; reasoning inductively; and exercising judgment.
>
> *Machines are better at:* responding quickly to control signals; applying great force smoothly and precisely; storing information briefly or erasing it completely; and reasoning deductively.

It is obvious that during the intervening half century some of Fitts's assertions no longer ring fully true. Energy detection, pattern recognition, and information storage and retrieval have made considerable progress, though inductive reasoning and judgment remain elusive. Sheridan (2000) lists the following problems of function allocation:

1. Computers, automation and robotics offer ever greater capability, but at the cost of greater system complexity and designer bewilderment, making the stakes for function allocation ever higher than before.
2. Proper function allocation differs by process stage (acquisition of information, processing and display of information, control action decision, execution of control).
3. Automation appears most promising at intermediate complexity, but the bounds of "intermediate" are undefined.
4. Human-centered design," while an appealing slogan, is fraught with inconsistencies in definition and generalizability.
5. "Naturalistic" decision-making and "ecological" design are sometimes incompatible with normative decision theory.
6. Function allocation IS design, and therefore extends beyond science.
7. Living with the technological imperative, letting our evolving machines show us what they can do, acceding or resisting as the evidence becomes clear, appears inevitable.

In spite of our best efforts to cope with these and other problems of function allocation, error and dispute over allocation criteria are human nature. Perhaps that is part of the Darwinian reality, the requisite variety, the progenitor of progress. At least we have it in our power to say no to new technology, or do we?

6.4 Levels of Automation

Sheridan and Verplank (1979) and Parasuraman et al. (2000) articulated a hierarchy of levels of automation, where it is up to the designer to decide which level is appropriate to the task. In increasing degree of automatic control:

1. The computer offers no assistance; the human must do it all.
2. The computer suggests alternative ways to do the task.
3. The computer selects one way to do the task, and:
4. ---executes that suggestion if the human approves, or
5. ---allows the human a restricted time to veto before automatic execution, or
6. ---executes automatically, then necessarily informs the human, or
7. ---executes automatically and informs the human only if asked.
8. The computer selects, executes, and ignores the human.

No one level guarantees system reliability or safety, and the different failure examples above can be said to have occurred at different levels. Multiple levels can be made available to the human user, as they exist now in the autopilot/flight management system of the commercial aircraft and might exist for the air traffic controller. A regression toward manual control is recommended for anomalous situations that cannot be handled by higher levels of automation.

6.5 Characteristic Biases of Human Decision-Makers

Research to date makes it clear that humans have difficulty with quantification, and systematically deviate from rational norms (such as Bayesian probability updating).

1. Decision makers are fairly good at estimating means, variances, and proportions—unless probability is close to 0 or to 1. Humans tend to regard very large numbers, such as 10^5, 10^6, 10^7, and 10^8, as all the same, even though they may be orders of magnitude different from one another. This may also be said for very small numbers, e.g., 10^{-5},10^{-6}, 10^{-7}, and 10^{-8}. Humans are much better at making ratio comparisons between numbers where the ratio is not greater than 1,000. Winkler and Murphy (1973) showed that weather forecasters are one of few groups who are good at quantitative prediction.
2. Decision makers do not give as much weight to past outcomes as Bayes' rule would indicate (Edwards, 1968). Probabilities of alternative hypotheses or propositions tend to be estimated much more conservatively than Bayes' theorem of probability updating would predict.
3. Decision makers often neglect base rates (Tversky and Kahneman, 1980; Edwards, 1968), a common tendency in which recent evidence is overweighted and previous evidence is neglected. Of course, for a rapidly changing situation (a non-stationary statistical process) this may be rational.

4. Decision makers tend to ignore the reliability of the evidence (Tversky and Kahneman, 1974).
5. Decision makers are not able to treat numbers properly as a function of whether events are mutually independent or dependent. They tend to overestimate the probability of interdependent events and underestimate the probability of independent events (Bar Hillel, 1973).
6. Decision makers tend to seek out confirming evidence and disregard disconfirming evidence (Einhorn et al., 1978).
7. Decision makers are overconfident in their predictions (Fischhoff, Slovic and Lichtenstein, 1977).
8. Decision makers tend to infer illusory causal relations (Tversky and Kahneman, 1973).
9. Decision makers tend to recall having had greater confidence in an outcome's occurrence or non-occurrence than they actually had before the fact (Fischhoff, 1975). This is called *hindsight bias*: "I knew it would happen."

Tversky and Kahneman (1974) showed that the above discrepancies can be attributed to three heuristics:

1. *Representativeness* or *framing.* The probability of an event belonging to a category B is judged by considering how representative A is of B. Since long series of heads or tails when flipping coins is considered unrepresentative, people are likely to predict the other event on the next trial. The illusion of validity occurs when people treat highly correlated events as though they are independent, thereby adding to the weight of one hypothesis. Judgments can be quite different depending on how the question of proposition is framed, even though the probabilities and consequences remain the same. For example, with medical interventions, the results are different depending on whether the outcomes are posed as gains or losses (Tversky and Kahneman, 1981). People will overestimate the chance of a highly desirable outcome and underestimate the chance of an undesirable one (Weber, 1994). A food that is "90 percent lean" is more acceptable than food that is 10 percent fat. Fifty percent survival is more acceptable than 50 percent mortality. In other words, a glass half full is more believable than a glass half empty.
2. *Availability.* An event is considered more likely if it is easy to remember, e.g., an airplane crash. Evans (1989) argues that people's consideration of available but irrelevant information is a major cause of bias.
3. *Anchoring and adjustment.* This is the idea that people update their degree of belief by making small adjustments (based on the new evidence) relative to the last degree of belief (Hogarth and Einhorn, 1992).

There are large individual differences between decision makers in several dimensions (Pew and Mavor, 1998):

1. *Degree of decision-making experience*. This includes sophistication in relation to concepts of probability, and how probability combines with level of consequences to form a construct of risk or safety.

2. *Acceptance of risk.* Most decision-makers are risk-averse, while some are risk-neutral or even risk-prone.
3. *Relative weighting on cost to self* in comparison to cost to some adversary or other party.
4. *Tendency to decide on impulse* in contrast to deliberation.

None of the above can be called irrational; they are simply differences in decision-making style.

6.6 Human Controller's Mental Model and/or Automatic Control "Model" of Process: Divergence from Reality

When the controller's internal model of the process (either the human controller's mental model or the software model in the automatic control system) diverges from the process state, erroneous control commands (based on the incorrect model) can lead to an accident. For example, (1) the software does not know that an aircraft is on the ground and raises the landing gear, or (2) it does not identify an object as friendly and shoots a missile at it, or (3) the pilot thinks the aircraft controls are in speed mode but the computer has changed the mode to open descent and the pilot issues inappropriate commands for that mode, or (4) the computer does not think the aircraft has landed and overrides the pilot's attempts to operate the braking system. There were corresponding examples of these events in the above failure reviews.

Experience suggests that serious accidents involve human unsafe acts performed when the system enters scenarios and conditions not understood by the operator(s). Therefore, in designing systems there is a need to prospectively identify both the conditions expected by operators (called the "base case") and possible deviations from that expected situation. In conjunction with such deviations it is then important to identify operational vulnerabilities that would precipitate human unsafe acts under the off-base conditions; examples include inadequate procedures, lack of operator knowledge, and operator biases. The nuclear safety industry has developed a procedure called ATHEANA to perform such an analysis (Forester et al., 2004). A newer approach to human reliability termed *reliance engineering* has emerged, emphasizing the organizational factors that lead to human misunderstanding and error (Hollnagel et al, 2006).

6.7 Undermonitoring: Over-reliance on Automation, and Trust

Undermonitoring of and overly trusting in automation have commonly been cited as contributors to human-automation failure. Trust of automation by humans is not a variable that engineers have customarily dealt with; their preference is reliability based on after-the-fact tabulations and statistical analysis of system and subsystem failures. However, trust of automation has recently become a salient consideration by human factors engineers (Lee and See, 2004; Gao and Lee, 2006). Two important questions are: how does trust grow or diminish as a function of automation performance, and under what circumstances do human operators tend to over- or undertrust the automation.

An explanation for undermonitoring of automation that complements the trust theory is one based on attention. Common sense tells us that there is nothing to be gained by attending to very reliable automation with low downside risk (e.g., home heating systems) until after failure is evident. Attending to imperfect automation is also diminished when the operator is engaged in other tasks that require attention. Studies of eye movements support this view. For example, Metzger and Parasuraman (2001) had observed air traffic controllers monitoring a radar display for separation conflicts while simultaneously accepting and handing off aircraft to and from their sector, managing electronic flight strips, and using data links to communicate with pilots. They were assisted by a conflict probe aid that graphically predicted the future courses (up to 8 minutes) of pairs of aircraft in the sector. The automation was highly reliable and reduced the time that controllers took to call out the conflict. However, in one scenario the automation did not point out the conflict because it did not have access to the pilot's intent to change course. Not surprisingly, controllers were either considerably delayed or missed the conflict entirely. Eye movement analysis showed that those controllers who did not detect the conflict had fewer fixations of the radar display compared to when they had been given the same conflict scenario without the conflict probe aid. This finding is consistent with the view that over-reliance on automation is associated with reduced attention allocation compared to manual conditions.

6.8 Mystification and Naive Trust

Human supervisors of computer-based systems sometimes become mystified and awed by the power of the computer, even seeing it as a kind of magical authority figure. This leads quite naturally to naive and misplaced trust. This was particularly well articulated by Norbert Wiener (1964), who used as a metaphor a classic in horror literature, W. W. Jacobs' *The Monkey's Paw*. The metaphor is salient.

> In this story, an English working family sits down to dinner. After dinner the son leaves to work at a factory, and the old parents listen to the tales of their guest, a sergeant major in the Indian army. He tells of Indian magic and shows them a dried monkey's paw, which, he says, is a talisman that has been endowed by an Indian holy man with the virtue of giving three wishes to each of three successive owners. This, he says, was to prove the folly of defying fate.
>
> He claims he does not know the first two wishes of the first owner, but only that the last was for death. He was the second owner, but his experiences were too terrible to relate. He is about to cast the paw on the coal fire when his host retrieves it, and despite all the sergeant major can do, wishes for £200.
>
> Shortly thereafter there is a knock at the door. A very solemn gentleman is there from the company that has employed his son and, as gently as he can, breaks the news that the son has been killed in an accident at the factory. Without recognizing any responsibility in the matter, the company offers its sympathy and £200 as solatium.

The theme here is the danger of trusting the magic of the computer when its operation is singularly literal. If you ask for £200 and do not express the condition that you do not wish it at the cost of the life of your son, you will receive £200 whether your son lives or dies.

To a naive user the computer can be simultaneously so wonderful and intimidating as to seem faultless. If the computer produces other than what its user expects, that can be attributed to its superior wisdom. Such discrepancies are usually harmless, but if they are allowed to continue they can, in some complex and highly interconnected systems, endanger lives. As new computer and control technology is introduced, it is crucial that it is accepted by users for what it is—a tool meant to serve and be controlled ultimately by human beings (Sheridan, 1992). The story of the monkey's paw, highlighted by Wiener, the "father of cybernetics" in his last (Pulitzer Prize-winning) book, is a lesson about the hubris of technology that is relevant to planning NGATS.

6.9 Remedies for Human Error

Given some understanding of the error situation, the usual wisdom for keeping so-called bad errors in check, according to the human factors profession, is (in order of efficacy) (Sheridan, 2002):

5. *Design to prevent error.* Provide immediate and clear feedback from an inner loop early in the consequence chain. Provide special computer aids and integrative displays showing which parts of the system are in what state of health. Pay attention to cultural stereotypes of the target population. For instance, since the expectation in Europe is that flipping a wall switch down turns a light on, when designing for Europeans, do not use the American stereotype of flipping the wall switch up to turn the light on. Use redundancy in the information, and sometimes have two or more actors in parallel (although this does not always work). Design the system to forgive and to be "fail safe" or at least "fail soft" (i.e., with minor cost).
6. *Train operators.* Train operators about the mental models appropriate for their tasks, and make sure their mental models are not incorrect. Train operators to admit to and think about error possibilities and error-causative factors; even though people tend to catch errors of action, they tend not to catch errors of cognition. Train operators to cope with emergencies they have not seen before, using simulators where available. Use skill maintenance for critical behaviors that need to be exercised.
7. *Restrict exposure to opportunities for error.* To avoid inadvertent actuation, ensure the fire alarms or the airplane exit doors have two or more activation steps and use key locks for certain critical controls that are seldom required. However, be aware that this limits the operator's opportunity for access in crisis.
8. *Alarm or warn.* Too many alarms or warnings on the control panels or in printed instructions tend to overload or distract the observers so they become conditioned to ignore them. Tort lawyers would have everyone believe that warnings are the most essential means to ensure safety. They may be the best way to guard against lawsuits, but they are probably the least effective means to achieve safety from a human factors viewpoint.
9. *Consider which behaviors are acceptable, which errors are likely, and what to do about them.* It is better to design robust systems that tolerate human variability than expect people to be error-free zombie automatons. If automation is indicated, keep the operator knowledgeable about what the automation procedures allow humans to take over if the automation fails, and engender some responsibility for doing this. Do not be too quick to blame the operators closest to the apparent error occurrence. Tilt toward blaming the system, and be willing to look for latent errors.

6.10 Can Behavioral Science Provide Design Requirements to Engineers?

The engineering of hardware and software has become very sophisticated. Design data and mathematical modeling tools abound, backed up by well-established laws of physics. Human understanding of human behavior is much less developed. The applied discipline of human factors engineering (or human-machine systems), like the discipline of medicine, is mostly based on empirical study, with relatively few equations or substantially “hard” laws. A tendency of design engineers has been to dismiss human factors for this reason, or to begrudgingly accept design reviews by human factors professionals late in the system design cycle. But this has often proven ineffective because at this point the human factors professionals can do little beyond raising problems and are seen as naysayers who are in opposition to the proponents of the almost completed system designs.

Providing design requirements that are directly usable by design engineers is the challenge for human-automation interaction and for human factors engineering in general. Human performance in defined tasks must become representable in the same terms as those used by engineers—in both static and fast-time dynamic simulations that include mathematical models of human operators as well as other system components. Real-time simulations with real humans in the loop can lead the way.

6.11 The Blame Game: The Need to Evolve a Safety Culture

The current ATM culture supports what has been called a “blame game;” all failures, including infractions of safety rules, have causes. Responsibility for these failures must be determined and penalties meted out. This approach to safety is exacerbated by the decades-old standoff between labor and management within ATC operating staff. One result is that infractions are only partially reported; line controllers are loath to call attention to their own or to their colleagues’ shortcomings.

A different, and many believe a more enlightened approach is to have an operating culture acknowledge that errors will happen—one where operating staff are encouraged not only to report but also to suggest ways to ameliorate the factors that allow the errors to occur. The American statistician/industrial engineer W. Edwards Deming contributed greatly to U.S. military production during World War II and lectured extensively in Japan after the war. He taught the Japanese quality-control techniques and about the importance of worker sensitivity to their own work efficiency. He also fostered open communication about errors and problems both horizontally among worker groups and vertically between layers of management. The techniques worked. The Japanese became the global model for industrial production and Deming became a demigod in Japan.

More recently the Institute of Medicine of the U.S. National Academy of Sciences (Kohn et al., 2000) published the report *To Err Is Human*, calling upon the medical community to desist from their well known “blame game” in which medical errors are closely guarded and underreported. Physicians operate in fear of malpractice suits. Safety performance data are not shared among hospitals, and physician training emphasizes personal responsibility but not teamwork or systems improvement thinking (along the lines of Deming). Largely as a result of this report there are new efforts to change the culture. No one is saying it will be easy. The U.S. culture of litigation also needs to be changed. One Harvard medical malpractice attorney told the writer that in her experience when physicians being sued openly admitted their errors, juries were always understanding and the defendants were almost always acquitted.

NGATS may offer an opportunity to bring about a more enlightened safety culture in aviation.

6.12 Concluding Comment

The famous physicist Richard Feynman, in his last book *What Do You Care What Other People Think?* is quoted by Degani (2004) as describing the inspiration he received from a Buddhist monk who told him “To every man is given the key to the gates of heaven; the same key opens the gates of hell.” We can all agree with Degani when he concludes, “ I believe the same applies when it comes to designing and applying automation.” Automation may be a key to a much improved air transportation system, but it can also precipitate disaster.

7.0 REFERENCES

Air Safety Week. Human factors issues emerge from Concorde crash investigation, Feb. 11, 2002.

Bainbridge, L .(1987). The ironies of automation. In *New Technology and Human Error*, eds. J. Rasmussen, K. Duncan, and J. Leplat. London: Wiley.

Bar Hillel, M. (1973). On the subjective probability of compound events. *Organizational Behavior and Human Performance* 9:396-406.

Billings, C. E. (1997). *Aviation automation: The search for a human centered approach.* Mahwah, NJ: Erlbaum.

Casey, S. (2006). *The Atomic Chef, and Other True Tales of Design, Technology and Human Error*. Santa Barbara CA: Aegean.

Decker, S., and Hollnagel, E. (1999). *Coping with Computers in the Cockpit*. Brookfield, VT: Ashgate.

Degani, A. (2004). *Taming HAL: Designing Interfaces Beyond 2001.* New York: Palgrave MacMillan.

Edwards, W. (1968). Conservatism in human information processing. In *Formal Representation of Human Judgment, ed.* B. Kleinmutz, 17-52. New York: Wiley.

Einhorn, H.J., Hogarth, R.M., and Robin, M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review* 85:395-416.

Evans, J.B.T. (1989). *Bias in Human Reasoning: Causes and Consequences*. Mahwah, NJ: Erlbaum.

Fischhoff, B. (1975). Hindsight = foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance* 1:288-299.

Fischhoff, B., Slovic, P., and Lichtenstein, S. (1977). Knowing with certainty: the appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance* 3:552-564.

Fitts, P.M. (1951). Human engineering for an effective air navigation and traffic control system. Ohio State University Foundation report. Columbus, OH.

Forester, J., Bley, D., Cooper, S., Lois, E., Siu, N., Kolaczkowski, A., and Wreathall, J. (2004). Expert elicitation approach for performing ATHEANA quantification. *Reliability Engineering and System Safety* 83:207-220.

Funk, K., Lyall, B., Wilson, J., Vint, R., Miemcyzyk, M., and Suroteguh, C. (1999). Flight deck automation issues. *International Journal of Aviation Psychology 9:*125–138.

Gao, J., and Lee, J.D. (2006). Extending the decision field theory to model operators' reliance on automation in supervisor control situations. *IEEE Transactions on Systems, Man, and Cybernetics*, Part A, 36(5):943-959.

Hogarth, R.M., and Einhorn, H.J. (1992). Order effects in belief updating: the belief adjustment model. *Cognitive Psychology* 25:1-55.

Hollnagel, E., Woods, D., and Leveson, N. (2006). *Resilience Engineering.* Williston, VT: Ashgate.

Infield, S., and Corker, K. (1997). The culture of control: free flight, automation and culture. In *Human-Automation Interaction: Research and Practice*, eds. M. Mouloua and J. Koonce. Lawrence Erlbaum Associates, 279-285.

Kletz, T.A. (1982). Human problems with computer control. *Plant/Operations Progress*, 1(4), October.

Kohn, L.T., Corrigan, J.M., and Donaldson, M.S. (2000). *To Err Is Human*. Washington, DC: National Academy Press.

Lee, J., and See, J. (2004). Trust in automation: designing for appropriate reliance. *Human Factors 46:*50–80.

Leplat, J. 1987. Occupational accident research and systems approach. In *New Technology and Human Error*, eds. Rasmussen, J., Duncan, K., and Leplat, J., 181–191, New York: Wiley.

Leveson, N.G. (2001). Evaluating accident models using recent aerospace accidents. Technical Report, MIT Dept. of Aeronautics and Astronautics. Available at http://sunnyday.mit.edu/accidents.

Leveson, N.G. (2004). A new accident model for engineering safer systems. *Safety Science*, 42(4):237-270.

Leveson, N.G., Allen, P., and Storey, M.A. (2002). The analysis of a friendly fire accident: using a systems model of accidents. *Proceedings of the 20th International Conference on System Safety.*

Metzger, U., and Parasuraman, R. (2001). Automation-related "complacency": Theory, empirical data, and design implications. In *Proceedings of the Human Factors and Ergonomics Society 45th Annual Meeting*, 463–467. Santa Monica, CA: Human Factors and Ergonomics Society.

Michaels, D., and Pasztor, A. As programs grow complex, bugs are hard to detect; a jet's roller coaster ride. *Wall Street Journal,* May 30, 2006.

National Transportation Safety Board (1973). *Eastern Air Lines, Inc., L-1011, N310EA, Miami, Florida, December 29, 1972* (AAR-73-14). Washington, DC.

National Transportation Safety Board. (1998a). *Brief of accident NYC98FA020.* Washington, DC.

National Transportation Safety Board. (1998b). *Safety recommendation letter A-98-3 through -5, January 21,1998.* Washington, DC.

Norman, D. A. (1990). The problem with automation: inappropriate feedback and interaction, not "overautomation." *Philosophical Transactions of the Royal Society* (London), B237:585–593.

Nunes, A., and Laursen, T. (2004). Identifying the factors that contributed to the Ueberlingen midair collision. *Proc. Annual Meeting Human Factors and Ergonomics Society*, New Orleans, Sept. 2004.

Parasuraman, R., Sheridan, T.B., and Wickens, C.D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man and Cybernetics*, SMC 30(3):286-297.

Perrow, C. (1984). *Normal Accidents: Living with High Risk Technologies.* NY: Basic Books.

Pew, R., and Mavor, A. (1998). *Modeling Human and Organizational Behavior*. Washington, DC: National Academy Press.

Reason, J. (1990). *Human Error*. Cambridge University Press, 1990.

Sarter, N.B., and Amalberti, R. (2000). *Cognitive Engineering in the Aviation Domain*. Mahwah, NJ: Erlbaum.

Sarter, N., and Woods, D. D. (1995). How in the world did we ever get into that mode? Mode error and awareness in supervisory control. *Human Factors 37:*5–19.

Sheridan, T. B. (1992). *Telerobotics, Automation and Human Supervisory Control*. Cambridge, MA: MIT Press.

Sheridan, T.B. (2000). Function allocation: algorithm, alchemy or apostasy? *International Journal of Human-Computer Studies* 52:203-216.

Sheridan, T.B. (2002). *Humans and Automation*. New York, NY: Wiley.

Sheridan, T.B., and Verplank. W.L. (1979). Human and computer control of undersea teleoperators. *Man-Machine Systems Laboratory Report.* Cambridge, MA: MIT.

Sheridan, T.B., and Parasuraman, R. (2006). Human-automation interaction. In *Reviews of Human Factors and Ergonomics*, ed. R. Nickerson. Santa Monica: Human Factors and Ergonomics Society.

Sherry, L., and Polson, P. G. (1999). Shared models of flight management system vertical guidance. *International Journal of Aviation Psychology 9:*139–153.

Tversky, A., and Kahneman, D. (1973). Availability, a heuristic for judging frequency and probability. *Cognitive Psychology* 5:207-232.

Tversky, A., and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science* 185:1124-1131.

Tversky, A., and Kahneman, D. (1980). *Causal Schemes in Judgments Under Uncertainty*. New York: Cambridge University Press.

Tversky, A., and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science* 211:453-458.

Vicente, K. (2004). *The Human Factor.* New York: Routledge.

Weber, E. (1994). From subjective probabilities to decision weights: the effect of asymmetric loss functions on the evaluation of uncertain events. *Psychological Bulletin* 115:228-242.

Wickens, C.D., Mavor, A.S., Parasuraman, R., and McGee, J.P. (Eds.) (1988). *The Future of Air Traffic Control: Human Operators and Automation.* Washington, DC: National Academy Press.

Wiener, E.L., and Nagel, D.C. (1988). *Human Factors in Aviation*. New York: Academic Press.

Wiener, Norbert. *God and Golem, Inc.* Cambridge, MA: MIT Press, 1964.

Winkler, R.L., and Murphy, A.H. (1973). Experiments in the laboratory and in the real world. *Organizational Behavior and Human Performance* 10:252-270.

8.0 ACKNOWLEDGMENTS

The writer especially acknowledges the contributions of Prof. Kevin Corker of San Jose State University, who served as a valuable consultant throughout this project, and of Dr. Richard John, former director of the Volpe Center, who was instrumental in initiating the project and eliciting the author’s participation as principal investigator.

Several colleagues are also to be acknowledged for their pioneering research on human-automation interaction and human error and their reports on various accident situations. I particularly drew on the work of Prof. James Reason of Manchester University in the UK, Dr. Asaf Degani of NASA Ames Research Center, Prof. Raja Parasuraman of George Mason University, Dr. Steven Casey of Ergonomic Systems Design, Prof. Nancy Leveson of the Massachusetts Institute of Technology, and Prof. Kim Vicente of the University of Toronto.